

METHODS AND APPARATUS FOR RAPID CAPTURE OF PROGRAM IDENTIFIER DATA IN A BROADBAND TRANSCODER MULTIPLEXER

This application claims the benefit of U.S. Provisional Application No. 60/398,416 filed July 24, 2002.

FIELD OF THE INVENTION

The present invention relates to the field of digital multimedia. More specifically, the
5 present invention relates to the collection of digital multimedia program information from a plurality of MPEG transport streams.

BACKGROUND OF THE INVENTION

In a broadband multimedia distribution network, a transcoder/multiplexer (TMX) is often used to perform a number of different operations on the multimedia streams. Some of
10 these operations include basic multiplexing of the incoming MPEG-2 transport packets, adding or dropping streams that appear in the output stream, video bit rate transcoding, statistical multiplexing, MPEG-2 video splicing, multimedia program insertion, IP data encapsulation and many others. In order to perform these tasks, each multimedia program must be readily identified and distinguished from other programs in the transport streams.

15 A multimedia program may contain a variety of audio, video, animation, text, data, and other information. An MPEG-2 multimedia program, for example, may contain three elementary streams: a video, an audio and a data stream. In some instances, such a multimedia program may contain several additional audio or data streams, as well. In order to transmit or store such programs, the elementary streams are multiplexed to form the transport stream. The
20 transport stream structure is based upon packets of constant size (i.e., 188 bytes in MPEG-2 transport stream); each transport stream packet is comprised of a header and a payload. Although it is possible for a transport stream to contain only a single program stream, it is often required to transmit several multimedia programs simultaneously in a transport stream. In such cases, numerous audio, video, data and other elementary streams are combined
25 together, in a multiplexed fashion, to form the composite transport stream.

In order to facilitate identification and extraction of elementary streams, various program identification fields and mapping tables are dispersed throughout the transport stream. The information contained in these tables may be utilized to uniquely identify a specific multimedia program and/or elementary stream carried within a single or plurality of multiplexed transport streams. In a broadband multimedia distribution network, a receiver (e.g., a TMX) should be able to readily identify and extract any program stream and its constituent elementary streams from the plurality of transport streams at its disposal. In prior art systems, the program information (e.g., information contained in MPEG Section Tables) is collected sequentially. Sequential collection of such program information from a large number of MPEG transport streams could take several minutes. This is an unreasonable amount of time and would make a TMX receiving such a large number of transport streams practically unusable.

It would be advantageous to provide a method and apparatus for collecting the program information from a plurality of transport streams simultaneously and efficiently. It would also be advantageous to enable collection of the program information asynchronously and in parallel with the processing of the program information.

The methods and apparatus of the present invention provide the foregoing and other advantages.

SUMMARY OF THE INVENTION

The present invention relates to the collection of program information from a plurality of multimedia transport streams, such as MPEG transport streams. The program information may be in the form of Section Tables carried in the MPEG transport streams.

5 It is an object of the present invention to provide an apparatus and method for improving the collection rate of all program information contained in one or more multimedia transport streams.

The present invention enables collection of program information asynchronously and in parallel with the processing of the information. In other words, the present invention
10 enables the collection of program information on one port, such as port 2, while program information is being processed on a different port, such as port 1.

A further advantage of the present invention is that the program information may be collected as it appears in the transport stream. This aspect of the present invention does not affect the MPEG standard sequence of Section Table dependence. For example, the MPEG
15 standard requires that the Program Association Table (PAT) be collected before the Program Map Table (PMT), and the present invention does not change or affect that requirement. However, rather than collecting the PMT's in order of ascending Packet Identification (PID) Codes, the present invention will collect the PMT PID's in the order they appear in the transport stream.

20 Another feature of the present invention is the ability to organize, monitor, and match each request for the program information based on a plurality of identification information such as, but not limited to, the PID, Table ID, Table ID Extension, Version Number, Current/Next Indicator and Section Number. At any particular moment, there could be, e.g., anywhere from zero to 1000 or more outstanding requests to collect program information.
25 These requests, which may be initiated by different applications, are collected and matched against the program information as they are received from the transport stream.

In accordance with a further aspect of the present invention, the matching of collected and requested program identification information may be accomplished by sequentially searching the requested program information. Alternatively, more sophisticated search

algorithms, such as binary search algorithms, binary tree algorithms, or other type of search algorithms may be used to perform this task.

In accordance with the features of the present invention, all further inquiries regarding the status of requested program information may be blocked until all requested program
5 information is entirely collected and successfully matched against the received program information from the transport streams. Alternatively, the application that has initiated the request may periodically poll the status of that request, or a group of requests, in order to assess the availability of program information.

The present invention may also be used to collect Digital Queue Tone messages for
10 advertisement insertion and/or for digital program splicing in the multimedia transport stream.

Further, the present invention may be used to provide for the collection of Internet Data contained in the MPEG transport stream.

Those skilled in the art will appreciate that the methods and apparatus of the present invention may also be used to collect other types of data from a plurality of transport streams.
15 These and other objects, features and advantages of the present invention will become apparent to those skilled in the art upon reading of the following detailed description in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will hereinafter be described in conjunction with the appended drawing figures, wherein like numerals denote like elements, and:

Figure 1 shows an exemplary arrangement of MPEG-2 transport stream packets;

5 Figure 2 shows an exemplary arrangement of MPEG-2 Program Specific information (PSI);

Figure 3 shows a task model for an example embodiment of the invention;

Figure 4 is a block diagram of a transcoder/multiplexer (TMX) which can be implemented in accordance with the present invention;

10 Figure 5 is a system-level diagram of an example multimedia network in accordance with the invention;

Figure 6 is a block diagram of a network hub site having a TMX that can be implemented in accordance with the present invention;

15 Figure 7 is a block diagram of an example headend implementation in accordance with the invention; and

Figure 8 is a block diagram of example transmission side components in accordance with the invention.

DETAILED DESCRIPTION OF THE INVENTION

The ensuing detailed description provides preferred exemplary embodiments only, and is not intended to limit the scope, applicability, or configuration of the invention. Rather, the ensuing detailed description of the preferred exemplary embodiments will provide those skilled in the art with an enabling description for implementing a preferred embodiment of the invention. It should be understood that various changes may be made in the function and arrangement of elements without departing from the spirit and scope of the invention as set forth in the appended claims.

A multimedia transport stream, such as an MPEG transport stream, is often comprised of several program streams that are multiplexed to form a composite stream. Each multimedia program stream, in turn, may further contain several Elementary Streams ("ES"). These elementary streams may include audio, video, data and other types of information. In some instances it is possible for a multimedia program (e.g., a video program) to contain a plurality of sound tracks and/or other data streams such as subtitles and closed caption channels. Thus, a composite transport stream may be considered as a conglomeration of various elementary streams that are multiplexed and transported simultaneously.

An elementary stream is essentially an endless string of bits representing images, video, sound, etc. In order to facilitate the transmission, delivery and storage of such multimedia content, each elementary stream is broken up into discrete blocks or packets to form what is known as the Packetized Elementary Stream (PES). Each PES packet is comprised of a data section, containing the digitized multimedia data (e.g., compressed video, audio, text, etc.) as well as other fields that are used for synchronization, temporal alignment, stream identification and other purposes. These additional fields collectively comprise the header portion of a PES packet. The information contained in the PES header fields is subsequently used at the receiver to identify and extract the appropriate packets from the transport stream.

The composite transport stream is generally made up of many elementary streams from different programs. For example, an MPEG-2 transport stream used in a digital satellite TV broadcast system is generally comprised of several TV programs plus a variety of audio

and text streams per TV program. The transport stream is structurally arranged as blocks of information with a constant size. In the case of an MPEG-2 transport stream, these packets are 188 bytes long. The transport stream packets should not be confused with previously described PES packets, which are normally larger and may vary in size.

5 Fig. 1 illustrates the arrangement of MPEG-2 transport stream packets 10, as is well known in the prior art. Fig. 1 indicates that each 188-byte packet is comprised of a header 12 and a payload area 14. The header is further comprised of several other fields 16 that are used for identification, synchronization and other purposes. For example, the header begins with a unique synchronization field (sync byte) that is used to identify the start of each packet. The
10 header also contains a Packet Identification (PID) Code that uniquely identifies each elementary stream. A receiver interested in accessing an elementary stream can simply monitor the packets of an incoming transport stream and select only those packets with a desired PID value.

15 In a multiplexed transport stream, the presence of PID values alone may not be sufficient to uniquely identify all the components of a multimedia program. For example, a demultiplexer must not only know the PID values of specific elementary streams, but it also needs to know which audio stream is associated with which video stream. In order to overcome this ambiguity, additional Program Specific Information (PSI) is often included in the transport stream to clarify the association of different elementary streams with each other.

20 Fig. 2 illustrates an example of PSI fields and their usage in an MPEG-2 transport stream, as is known in the prior art. The demultiplexer, at start up, has no prior knowledge of composition of the transport stream. It must initially look for packets with PID of zero. This PID is reserved for the Program Association Table (PAT). The PAT contains a listing of all programs in the transport stream as well as the PID values for Program Map Tables (PMT)
25 associated with each multimedia program. For example, the PAT 20 of Fig. 2 includes three programs and three PID values corresponding to the PMT's of each program. The PMT 22 for program 1 (i.e., PMT with PID=35) further contains the PID values for one video, four audio and one data streams, all associated with program 1. Similarly, the PMT 24 for program 2 (i.e., PMT with PID=50) contains the PID values for one video, two audio and one data

streams. Program 0 and its corresponding PID (i.e., PID=18) identify the Network Identification Table (NIT) 26. The NIT contains information about other available transport streams. For example, in a satellite broadcast network, NIT 26 would contain information regarding orbital position, the polarization, carrier frequency and other information relating to the satellite.

The presence of PSI, such as shown in the exemplary PSI tables of Fig. 2, makes it possible to uniquely identify and decode any multimedia program, and its associated elementary streams in the presence of multiple programs and/or transport streams. Nevertheless, the task of searching a large number of transport streams with an even larger number of elementary streams can be very time consuming. This task may be further complicated if simultaneous requests for program information are waiting to be completed. It is an object of the present invention to provide more efficient methods and apparatus for collecting, organizing and matching such program information as the incoming packets from various transport streams are monitored. Using the methods and apparatus of the present invention, the collection of Section Tables from forty or more MPEG transport streams can be reduced from a few minutes to less than one second.

The present invention is able to organize, monitor, and match each request for program information based on parameters such as PID, Table ID, Table ID Extension, Version Number, Current/Next Indicator, and Section Number, etc. The requests are matched and collected as the program information is received from the transport stream.

Referring to Fig. 3, there is shown a high-level block diagram of a Task Model for an exemplary embodiment of the present invention. When a request for program information is initiated by an application, a structure containing all the relevant program information is generated. This request is made in the context of the Application's Task. The request structure contains all the information required to enable collection of the program information. By way of example, and not by way of limitation, some of the information that comprises the request structure includes the Port number, PID, PID Alias, Table ID, Table ID Extension, Current/Next Indicator, Section Number, and the like. The request structure also acts as an

interface between the Application Tasks 30, 32 and the Collection Task 34, as depicted in Fig. 3.

When the request for program information is enabled, this structure is moved to a linked list of active requests, i.e., the Active List (a linked list, as is well known in the prior art, is a special type of list that can be dynamically changed in size; dynamic structures can be changed in size and are unique in that they allow the insertion or deletion of any element in the list quickly and efficiently). There is an Active List associated with each port that accepts incoming transport stream packets. When an MPEG transport packet is received, an Interrupt Service Routine 36 determines which port the MPEG packet came from and passes this information to the Collection Task 34. The Collection Task then uses the port number and performs a search of the associated Active List. The search commences with matching of the received PID against the PID values on the Active List. If a match is found, the request is moved to the Process Request linked list. The mere matching of the PID, however, is not sufficient since there may exist multiple requests with the same PID, but different Section Table ID's, for example, contained in the request. Once the request is moved to the Process Request list, the remaining items of the request, such as the Table ID, Table ID Extension, Version Number and Section Number, and others, are compared against the corresponding information received from the transport packets. If a complete match between the received program information and items on the Process Request list is found, the request is moved to the Completed List and the requesting application is notified. If at any stage of the search, a mismatch is encountered, the request is moved back to the Active List, awaiting the next batch of received program information.

It is instructive to note that MPEG Section Tables may be contained in a single MPEG transport packet or may be split across several packets. An MPEG packet may contain a single Section Table or several. Thus, it may become necessary to conduct the search and compare procedures for several incoming packets multiple times for a single packet. It can be appreciated that many different search algorithms may be used to conduct the above described matching procedure between the incoming transport packet information and the contents of the Process Request lists. The simplest way is to sequentially examine each item on the

Process Request list in search of a match. This is also known as the linear search algorithm and while conceptually simple, this method can be computationally inefficient. Consequently other algorithms may be used to improve the search efficiency. For example, a binary search algorithm may be used where the original list is repeatedly divided in two halves and each half is searched until a match is obtained. Other well-known search methods such as Binary Tree Algorithms may be used, as well.

It can also be appreciated that using the features of the present invention, the collection and processing of the program information are effectively decoupled. This allows asynchronous collection and processing of the program information. For example, while the collection is taking place on port 1, the processing of the program information can take place on port 2. From the Application Task's point of view all requests for program information are filled in parallel. Once a complete match for a program request is fulfilled, the application may be notified so that decoding and/or further processing of the requested program(s) can take place. Using this methodology, the application can request to be pended until the requested program information is collected. Alternatively, the application may periodically poll the status of each request, or a group of requests, to assess the availability of program information. Accordingly, the Application Task can execute a Status Query, which returns the number of requests that are completed.

Referring to Fig. 4, there is shown a block diagram of a Transcoder/Multiplexer (TMX), which implements an example embodiment of the present invention. The Input Module (IM) 42 captures the transport stream signals that contain the MPEG video, audio, and Program Information tables. Each of these transport streams may contain a large number, for example anywhere from two to fifty different section tables. These may include a variety of tables, such as, but not limited to PAT's, PMT's, Program and System Information Protocol (PSIP) Tables, Event Information Tables (EIT's), etc. The Input Processor (IP) 44 captures and routes the PID's to the Transcoder/Multiplexer Card (TRC) 40 if the PID is video/audio, or to the Central Processing Unit (CPU) 46 if the PID contains a Program Information Table. The different section tables are identified and processed by the CPU, as

described above, in order to determine whether or not a match between the requested and received program information is obtained

Figures 5 to 7 are block diagrams showing possible systems where example embodiments of the present invention may be utilized. In Fig. 5, the TMX 50 which implements the present invention is used in a broadband environment with a variety of multimedia streams including satellite, Internet and terrestrial distribution and transmission networks. In Fig. 6, a TMX 60 is employed with multiple Asynchronous Serial Interface (ASI) inputs for multimedia streams; some of the input streams may have been originated remotely (e.g., carried over OC3) while others may have been locally produced as analog streams and fed into single channel video encoders 62 to produce compressed MPEG-2 streams. The TMX output is a multiplexed ASI output of multimedia streams suitable for digital distribution. In Fig. 7, the TMX 70 is utilized with multiple satellite dish networks as well as local encoders. This is an example of the TMX being used in a *statistical multiplex* mode with the present invention. In the implementation illustrated, there are four *statistical multiplex* groups each with twelve MPEG services. In this implementation, the present invention may be utilized to collect the PSI for forty-eight MPEG services simultaneously. It should be appreciated that implementations using a different number of statistical multiplex groups and/or a different number of services can be provided instead.

Figure 8 depicts an ATSC broadcast example in accordance with the invention in which a TMX 80 is used to multiplex data from several sources including an ATSC program encoder 82, network information (i.e., PSIP) generator 84, and local (e.g., disk-based) 86 and remote (e.g., Internet) 88 data server traffic. It is understood that the features of the present invention can be used in any environment that requires quick integration and processing of multimedia services such as networks with MPEG-transport distribution architecture. These services and operations may include, but are not limited to, basic packet multiplexing, statistical multiplexing, IP data encapsulation, video splicing, advertisement insertion, bit rate transcoding and generation of customized transport streams. All of the above services and operations require identification and extraction of program information from packetized MPEG transport streams and thus can benefit from the embodiments of the present invention.

When collecting Section Tables from such a large number of transport streams, the present invention improves the collection rate from several minutes to less than one second.

It should now be appreciated that the present invention provides advantageous methods and apparatus for collecting program information from a plurality of transport
5 streams simultaneously. Although the invention has been described in connection with various illustrated embodiments, numerous modifications and adaptations may be made thereto without departing from the spirit and scope of the invention as set forth in the claims.